

MobileSpeech: A Fast and High-Fidelity Framework for Mobile Zero-Shot Text-to-Speech

Shengpeng Ji^{*1,2}, Ziyue Jiang^{*1}, Hanting Wang¹, Jialong Zuo¹, Zhou Zhao^{†1,2}

Zhejiang University¹

Shanghai AI Laboratory²

shengpengji, zhaozhou@zju.edu.cn

Abstract

Zero-shot text-to-speech (TTS) has gained significant attention due to its powerful voice cloning capabilities, requiring only a few seconds of unseen speaker voice prompts. However, all previous work has been developed for cloud-based systems. Taking autoregressive models as an example, although these approaches achieve high-fidelity voice cloning, they fall short in terms of inference speed, model size, and robustness. Therefore, we propose MobileSpeech, which is a fast, lightweight, and robust zero-shot text-to-speech system based on mobile devices for the first time. Specifically: 1) leveraging discrete codec, we design a parallel speech mask decoder module called SMD, which incorporates hierarchical information from the speech codec and weight mechanisms across different codec layers during the generation process. Moreover, to bridge the gap between text and speech, we introduce a high-level probabilistic mask that simulates the progression of information flow from less to more during speech generation. 2) For speaker prompts, we extract fine-grained prompt duration from the prompt speech and incorporate text, prompt speech by cross attention in SMD. We demonstrate the effectiveness of MobileSpeech on multilingual datasets at different levels, achieving state-of-the-art results in terms of generating speed and speech quality. MobileSpeech achieves RTF of 0.09 on a single A100 GPU and we have successfully deployed MobileSpeech on mobile devices. Audio samples are available at <https://mobilespeech.github.io/>.

1 Introduction

In recent years, remarkable progress has been made in the development of text-to-speech (TTS) technology (Ren et al., 2019; Kim et al., 2021a; Huang et al., 2022a; Li et al., 2023). With advancements in

this field, current state-of-the-art TTS systems have demonstrated the ability to generate high-quality voices of unseen speakers while maintaining consistency in tone and rhythm (Chen et al., 2021; Casanova et al., 2022b). However, these systems are often trained on datasets of only a few hundred hours and often require fine-tuning with dozens of sentences during the inference stage, which significantly limits the generative capabilities of the models and is time-consuming and not user-friendly.

With the recent development of large-scale language models (Brown et al., 2020; Touvron et al., 2023) and the increase in training data for speech synthesis from hundreds of hours (Zen et al., 2019) to tens of thousands of hours (Kahn et al., 2020), current TTS models exhibit powerful zero-shot capabilities, enabling the cloning of unseen speakers' voices in just a few seconds (Wang et al., 2023; Zhang et al., 2023; Shen et al., 2023; Le et al., 2023; Jiang et al., 2023b; Kharitonov et al., 2023; Borsos et al., 2023b; Huang et al., 2023; Yang et al., 2023b). However, these models primarily focus on in-context abilities and fail to address issues related to inference speed and model deployment parameters. Additionally, all existing zero-shot TTS models are cloud-based and lack a mobile-based deployment approach. MobileSpeech is the first zero-shot TTS synthesis system that can be deployed on mobile devices. Firstly we think deployable mobile zero-shot TTS models aim to achieve the following goals:

- **Fast:** through our experiments, we have observed that the real-time factor (RTF) on a single mobile device often exceeds 8-10 times that of an A100 RTF for the same text. Therefore, the inference speed of the zero-shot TTS model must be significantly faster.
- **Lightweight:** to enable deployment on mobile or edge devices, the model size should be small, and the runtime memory footprint should be minimal.

^{*}Equal contribution.

[†]Corresponding author.

- **High similarity and diversity:** zero-shot TTS system should be able to clone timbre and prosody with just a few seconds of prompts, yielding diverse speech even when provided with identical text.
- **High quality:** to enhance the naturalness of synthesized speech, the model should pay attention to details such as frequency bins between adjacent harmonics and possess powerful duration modeling capabilities.
- **Robustness:** a highly robust TTS system is essential. A zero-shot TTS system should minimize occurrences of missing or repeated words.

To achieve these objectives, we have made our best efforts to reproduce several state-of-the-art generative models, which will be elaborated in the subsequent section on related work. We found current work is not applicable to mobile devices. Ultimately, based on the Natspeech framework (Ren et al., 2019, 2020) and the Mask generation module (Chang et al., 2022; Borsos et al., 2023b), we additionally design the SMD module and the speaker prompt component, taking into account the discrete codec architecture of speech. In summary, MobileSpeech contributes as follows:

- MobileSpeech is the first zero-shot TTS synthesis system that can be deployed on mobile devices. MobileSpeech achieves a good balance according to the above five evaluation metrics.
- In MobileSpeech, we additionally design the SMD module based on the hierarchical token structure of discrete speech codecs and design the Speaker Prompt module to maintain high similarity, high quality, low latency.
- Training MobileSpeech-english for 580 hours achieved SOTA inference speed and comparable audio quality. Training MobileSpeech-chinese for 40,000 hours resulted in SOTA performance in all aspects.
- MobileSpeech has been successfully deployed on mobile phones and is expected to be used by several hundred thousand users.

2 Related Work

2.1 Zero-shot TTS

Zero-shot speech synthesis refers to the ability to synthesize the voice of an unseen speaker based solely on a few seconds of audio prompt, also known as voice cloning. Due to its impressive performance, it has gained significant attention recently. Previous approaches in the field of voice cloning can be classified into speaker adaptation (Wang et al., 2018; Chen et al., 2021; Huang et al., 2022b) and speaker encoding methods (Casanova et al., 2022b; Arik et al., 2018; Kang et al., 2022). These models are often trained on smaller datasets and require a substantial amount of data for fine-tuning or employ highly complex encoders. These limitations greatly restrict the audio quality and further usage.

In recent months, with the advancement of generative large-scale models, a plethora of outstanding works have emerged. VALL-E (Wang et al., 2023) leverages discrete codec representations and combines autoregressive and non-autoregressive models in a cascaded manner, preserving the powerful contextual capabilities of language models. It can clone the voice of a target speaker with just a 3-second audio prompt. VALL-E X (Zhang et al., 2023) extends zero-shot TTS to multiple languages based on the cascaded structure of VALL-E. NaturalSpeech 2 (Shen et al., 2023) employs continuous vectors instead of discrete neural codec tokens and introduces in-context learning to a latent diffusion model. SpearTTS (Kharitonov et al., 2023) and Make-a-Voice (Huang et al., 2023) utilize semantic tokens to reduce the gap between text and acoustic features. VoiceBox (Le et al., 2023) is a non-autoregressive flow-matching model trained to infill speech, given audio context and text. Megatts (Jiang et al., 2023b,a), on the other hand, utilizes traditional mel-spectrograms, decoupling timbre and prosody and further modeling the prosody using an autoregressive approach.

However, none of the aforementioned works take into account the model’s speed and lightweight nature. It is well-known that autoregressive models often require large memory footprints and have slower iteration speeds, and diffusion models often require hundreds of sampling steps. Some work introduce a semantic token cascade structure, which will introduce more time overhead. However, compared to codecs, certain approaches based on mel-spectrogram exhibit lower generation diversity and

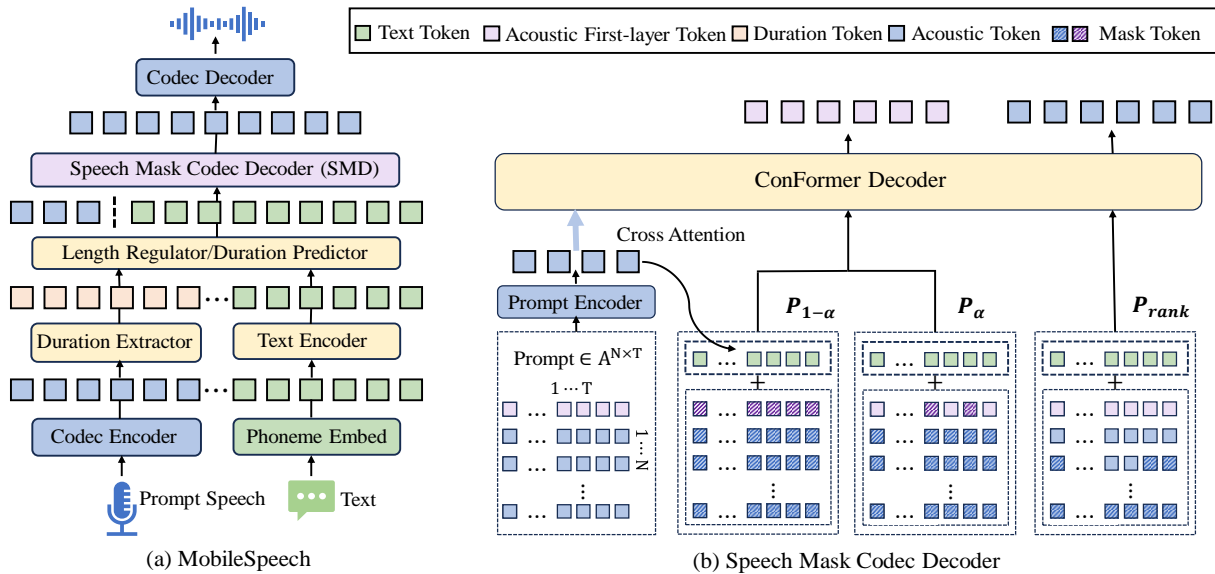


Figure 1: The overall architecture for MobileSpeech. In Figure (a) the Duration Extractor module is responsible for extracting prompt duration tokens from prompt acoustic tokens. SMD represents the generative module for target acoustic tokens. In Figure (b), we provide a detailed depiction of the multi-channel training process employed by the SMD module.

audio quality (Siuzdak, 2023). MobileSpeech, on the other hand, is the first zero-shot TTS system that can be deployed on mobile devices, addressing the fast speed, lightweight, high quality requirements.

2.2 Generative models

Generative models have been extensively utilized in various domains by employing techniques such as language models (Borsos et al., 2023a; Kreuk et al., 2022; Ji et al., 2023; Wang et al., 2023; Zhang et al., 2023), Variational Autoencoders (VAE) (Ren et al., 2021; Lee et al., 2022), Generative Adversarial Networks (GAN) (Kim et al., 2021b; Kong et al., 2020), Normalizing Flow (Kim et al., 2020; Miao et al., 2020), Mask Autoencoders (Chang et al., 2022) and diffusion models (Huang et al., 2022a). Considering the challenge of balancing perceptual quality and inference speed in zero-shot TTS, as well as the recent success of parallel decoding in text (Ghazvininejad et al., 2019), image (Chang et al., 2022), audio (Borsos et al., 2023b), and video (Villegas et al., 2022) generation tasks, MobileSpeech has adopted masked parallel generation as the audio synthesis approach. Furthermore, MobileSpeech has additionally designed the Speech Codec Mask Decoder (SMD) module and Speaker prompt based on the unique characteristics of acoustic tokens and the specific requirements of the zero-shot TTS task, aiming to address the afore-

mentioned trade-off. In Appendix A, we present the detail of discrete acoustic codec and in Appendix B we present the distinctions between mask-based generative models and other generative models in modeling discrete codecs.

3 MobileSpeech

In this section, we will first introduce the overall structure of MobileSpeech, followed by a detailed focus on the Speech Mask Codec Decoder module and the Speaker Prompt module. Furthermore, we will proceed to elaborate on the specific intricacies of MobileSpeech’s training and inference processes, with the explicit details of the model being presented in Appendix C.

3.1 Overall

As shown in Figure 1, MobileSpeech employs FastSpeech 2 as our baseline. The input text is processed to extract phoneme sequence, which are then passed through a text encoder to obtain the corresponding target text representation. Simultaneously, the input consists of speech prompts based on a discrete codec (Défossez et al., 2022). The prompt encoder compresses the channel dimension of the prompt codec acoustic tokens to one dimension. We specifically extract fine-grained prompt duration tokens from the prompt codec acoustic tokens and combine them with the target text, which are then fed into the duration prediction module and

length regularizer to obtain an expanded target text. The Speech Mask Codec Decoder serves as the core module. During the training process, we randomly select a portion of the aligned target text as the prompt to guide the generation of the remaining text. Based on the masking and parallel generation using RVQ features of the speech codec, we design a multi-channel training approach, priority training, and a mechanism that relies solely on text to generate acoustic tokens. Finally, the acoustic tokens are transformed into the final speech output using a pre-trained codec Tokenizer (Siuzdak, 2023).

3.2 Speech Codec Mask Decoder

Based on paired text-speech data $\{X, Y\}$, where $X = \{x_1, x_2, x_3, \dots, x_T\}$ represents the aligned text and Y denotes the representation of speech through a pre-trained discrete codec encoder (Défossez et al., 2022), namely $Y = C_{1:T,1:N} \in \mathbb{R}^{T \times N}$, we consider T as the downsampled utterance length, which is equal to the length of the text. N represents the number of channels for every frame. The row vector of each acoustic code matrix $C_{t,:}$ represents the eight codes for frame t , and the column vector of each acoustic code matrix $C_{:,j}$ represents the code sequence from the j -th codebook, where $j \in \{1, 2, \dots, N\}$.

Notably, in contrast to models such as VALL-E (Wang et al., 2023), it is acknowledged that obtaining prompt text in real-world scenarios is challenging, and to mitigate the impact of erroneous prompt text on the model. We have devised a text-agnostic speech prompt to guide the generation of target speech in MobileSpeech. Specifically, for a sequence with a time length of T , we randomly extract a segment of length k at the beginning to simulate prompts of arbitrary lengths, where the value of k falls within the range of:

$$\frac{T}{3} \leq k \leq \frac{2 \times T}{3} \quad (1)$$

By selecting the split point k , we obtain the corresponding prompt speech Y_{prompt} and the target text Y_{prompt} :

$$Y_{prompt} = C_{1:k,1:N}, X_{target} = X_{k:T} \quad (2)$$

Therefore, disregarding the codec channel count, the optimization objective of the Speech Codec Mask Decoder (SMD) module is to maximize the conditional probability function shown below:

$$P(C_{k:T,1:N} | C_{1:k,1:N}, X_{k:T}; \theta) \quad (3)$$

It is well known that the use of RVQ-based speech discretization tokens (Défossez et al., 2022; Siuzdak, 2023) induces a hierarchical token structure, where tokens from finer RVQ levels contribute less to the perceptual quality. This allows for efficient factorization and approximation of the joint distribution of the token sequence. We believe that considering this structural aspect holds the most promise for future advancements in long-sequence audio modeling. Therefore, in MobileSpeech, as depicted in Figure 1 (b), we divide the generation process into two parts: the generation of the first channel and the generation of the remaining channels at the same time, denoted as:

$$\begin{aligned} & P(C_{k:T,1:N} | C_{1:k,1:N}, X_{k:T}; \theta) \\ &= P(C_{k:T,1} | C_{1:k,1:N}, X_{k:T}; \theta) \\ & \times \prod_{j=2}^N P(C_{k:T,j} | C_{1:k,1:N}, X_{k:T}; \theta) \end{aligned} \quad (4)$$

For the generation of the first channel $P(C_{k:T,1} | C_{1:k,1:N}, X_{k:T}; \theta)$, we employ a mask-based generative model as our parallel decoder. We sample the mask $M_i \in \{0, 1\}^T$ according to a cosine schedule (Chang et al., 2022) for codec level i , specifically sampling the masking ratio $p = \cos(u)$ where $u \sim \mathcal{U}[0, \frac{\pi}{2}]$. and the mask $M_i \sim \text{Bernoulli}(p)$. Here, M_i represents the portion to be masked in the i -th level, while \bar{M}_i denotes the unmasked portion in the i -th level. The prediction of this portion is refined based on the prompt speech and the concat of target text and the unmasked portion of the first channel. Therefore, the prediction for this part can be specified as follows:

$$P(M_1 C_{k:T,1} | C_{1:k,1:N}, X_{k:T}, \bar{M}_1 C_{k:T,1}; \theta) \quad (5)$$

Due to the fact that in the inference process, we directly obtain the first-layer discrete tokens from the text, it often requires multiple iterations to iteratively select higher-confidence tokens to simulate the masked environment during training. In order to reduce the number of iterations, as shown in Figure 1 (b), MobileSpeech, during the training process, employs a certain probability to completely mask the acoustic tokens of the first channel. Additionally, completely masking the acoustic tokens of the first channel at different stages of model training can guide the model to generate from easier to harder instances, thus increasing the upper limit of the model. We control the two different ways of

generating the first channel through a probability parameter α , as shown below:

$$\begin{aligned} & P(C_{k:T,1} | C_{1:k,1:N}, X_{k:T}; \theta) \\ = & \alpha P(M_1 C_{k:T,1} | C_{1:k,1:N}, X_{k:T}, \bar{M}_1 C_{k:T,1}; \theta) \\ & + (1 - \alpha) P(M_1 C_{k:T,1} | C_{1:k,1:N}, X_{k:T}; \theta) \end{aligned} \quad (6)$$

For the generation of the remaining channels, similar to other models that randomly select a channel for training during the training process, we incorporate a weighting mechanism in the random selection process. We believe that the importance of channels decreases layer by layer. Therefore, we introduce the P_{rank} mechanism in Appendix D. When predicting the acoustic token of the j -th channel, we utilize the unmasked portion of the j -th layer \bar{M}_j and all acoustic tokens preceding the j -th layer for prediction, as shown below:

$$P(M_j C_{k:T,j} | C_{1:k,1:N}, X_{k:T}, C_{k:T,<j}, \bar{M}_j C_{k:T,j}; \theta) \quad (7)$$

3.3 Speaker Prompt

MobileSpeech, based on FastSpeech2, adopts an Encoder-Decoder architecture. Different prompt strategies were designed in both stages to encourage mobileSpeech to capture the prosodic and time-bre information of the prompt speech. During the experimental process, we observed that MobileSpeech generated speech with relatively uniform prosody and the distribution of durations exhibited excessive uniformity. Therefore, we abandoned the coarse-grained speech prompt guidance before the duration predictor. As illustrated in Figure 2, the Duration Extractor incorporates a Q-K-V attention layer to integrate different information. It first extracts fine-grained prompt duration values from the prompt acoustic token and feeds them into the prompt duration encoder along with positional information. Subsequently, the prompt duration is used to guide the prediction of the target text’s duration. Subsequent ablation experiments have confirmed that this approach can improve the occurrence of disfluent speech phenomena. Both the Duration Extractor and the Duration Predictor modules follow the same structure. The text is first passed through a normalization layer and then mapped to a Query. The Query, along with the prompt acoustic or prompt duration, undergoes a Q-V-K attention layer and a convolutional layer to obtain latent features. The latent features are then

added to the text through residual connections to yield the corresponding output results.

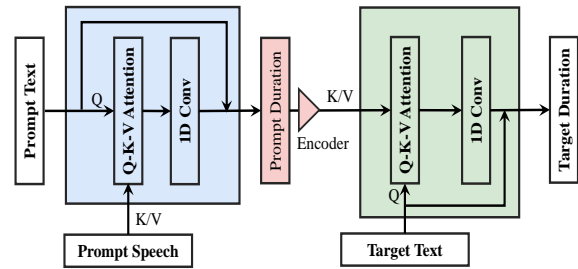


Figure 2: The process of obtaining target duration tokens from target text tokens and prompt speech tokens is depicted in the following manner: the blue module represents the Prompt Duration Extractor, while the green module represents the Duration Predictor.

In the Decoder stage of MobileSpeech, as described in Equation 1 and 2, during the training process, we randomly select a sequence of length k from the target text and speech to form a text-independent speech prompt. It is worth noting that, to ensure the robustness of the model, the speaker prompt in the Encoder and the speaker prompt in the Decoder are not identical sequences. In the Speech Mask Codec Decoder (SMD), all tokens of the input prompt acoustic in SMD are unmasked, and inspired by (Du et al., 2023), we do not include corresponding positional information and prompt encoder in order to preserve the original features of the acoustic tokens as one of the inputs to the SMD module. This is explicitly highlighted with a blue line in Figure 1 (b). Additionally, in the SMD module, we pass the acoustic tokens through the prompt encoder and apply cross-attention to allow the target text, as a weak condition token, to incorporate some information from the prompt acoustic tokens, thereby maintaining contextual capabilities.

3.4 Train and Inference

During the training process, the Duration Predictor and Duration Extractor are optimized using the mean square error (MSE) loss, where the extracted duration serves as the training target. We denote the losses at Duration Predictor and Duration Extractor as \mathcal{L}_{dur} and $\mathcal{L}_{promptdur}$, respectively. To improve alignment accuracy and reduce the information gap between the model input and output, we employ the Montreal forced alignment (MFA) (McAuliffe et al., 2017) tool to extract phoneme durations. The duration values are predicted in the logarithmic domain, which promotes a more Gaus-

sian distribution and facilitates training. The SMD module is optimized using the cross-entropy loss function, considering both the first channel and the discrete acoustic tokens randomly selected under the Prank strategy. We use \mathcal{L}_{smd}^1 and \mathcal{L}_{smd}^j to represent the losses of the first channel and the j -th channel, respectively. Where $2 \leq j \leq N$. Thus, the overall loss \mathcal{L} can be expressed as follows:

$$\mathcal{L} = \mathcal{L}_{smd}^1 + \mathcal{L}_{smd}^j + \mathcal{L}_{promptdur} + \mathcal{L}_{dur} \quad (8)$$

During the inference stage, we iteratively generate discrete acoustic tokens by incorporating the prompt speech into the original text state. For the generation of the first channel of acoustic tokens, we employ the confidence-based sampling scheme proposed by (Chang et al., 2022; Borsos et al., 2023b). Specifically, we perform multiple forward passes, and at each iteration i , we sample candidates for the masked positions. We then retain P_i candidates based on their confidence scores, where P_i follows a cosine schedule. As for the generation of the remaining channels, we directly select the token with the highest probability using a greedy strategy.

4 Experiments

4.1 Experiment Setup

Datasets We trained MobileSpeech on datasets of varying scales and languages, and evaluated its performance on datasets comprising speech prompts of different lengths and environments. This allowed us to measure the effectiveness of MobileSpeech. For the English section, we trained MobileSpeech exclusively on the libriTTS (Zen et al., 2019), which comprises 580 hours of data from 2,456 speakers. For the Mandarin Chinese section, MobileSpeech was trained on an internal dataset totaling 40,000 hours.

Baselines We consider the latest SOTA models as our baselines. For the English segment, we utilize YourTTS (Casanova et al., 2022b)¹ in its official version. We have successfully reproduced results close to the official ones for VALL-E (Wang et al., 2023), surpassing the reproduced results found online². As for NaturalSpeech2 (Shen et al., 2023), we employ the reproduced version available online. Additionally, we have achieved results similar to the official ones for Megatts (Jiang et al., 2023b). For the Chinese segment, we compare our

results with the open-source interface³ of the current state-of-the-art Chinese model Megatts2 (the version of the interface is November 2023) (Jiang et al., 2023a).

Training and Inference Settings MobileSpeech was trained on 40,000 hours of Mandarin Chinese data using 32 NVIDIA A100 40G GPUs. The discrete codec was trained end-to-end on the training dataset, with the encoder part of Encodec (Défossez et al., 2022) and the decoder part of Vocos (Siuzdak, 2023) being optimized through engineering modifications. For the 580-hour English training data, the discrete codec employed a pre-trained version. MobileSpeech was trained for 12 epochs on 4 NVIDIA A100 40G GPUs, with each batch accommodating 3500 frames of the discrete codec. We optimized the models using the AdamW optimizer with parameters $\beta_1 = 0.9$ and $\beta_2 = 0.95$. The learning rate was warmed up for the first 5k updates, reaching a peak of 5×10^{-4} , and then linearly decayed.

Automatic metrics We have aligned all our objective experiments with VALL-E (Wang et al., 2023). To evaluate speaker similarity (SPK) between the original prompt and synthesized speech, we employ WavLM-TDNN (Chen et al., 2022). However, due to updates in the repository, we have updated the feature extractor, but all our models have been tested by using the same metrics. For assessing automatic speech recognition (ASR) performance, we conduct ASR on the generated audio and calculate the word error rate (WER) compared to the original transcriptions. In this experiment, we utilize the HuBERT-Large (Hsu et al., 2021) model fine-tuned on LibriSpeech 960h as the ASR model. This model is a CTC-based model without language model fusion. To compute the real-time factor, we measure the ratio of the total length of generated speech to the total time taken by the model to generate the corresponding speech features. The inference time in the vocoder component will not be taken into account. The calculation is performed on an empty A100 GPU.

Human evaluation We conduct the MOS (mean opinion score) and CMOS (comparative mean opinion score) evaluation on the test set to measure the audio naturalness via crowdsourcing. We keep the text content and prompt speech consistent among different models to exclude other interference fac-

¹<https://github.com/Edresson/YourTTS>

²<https://github.com/enhuiz/vall-e>

³<https://www.volcengine.com/product/voicecloning>

Model	Data	WER ↓	SPK ↑	RTF ↓	MOS-Q ↑	MOS-P ↑	MOS-S ↑
GT Codec	-	2.4	0.871	-	4.41±0.08	4.28±0.10	4.45±0.06
YourTTS	640	7.7	0.504	0.22	3.71±0.11	3.59±0.14	3.68±0.12
VALL-E	60000	5.9	0.751	0.94	4.22±0.06	4.09±0.13	4.16±0.07
VALL-E-Continue	60000	3.8	0.734	0.94	4.12±0.12	4.13±0.10	4.11±0.11
NaturalSpeech2-Continue	580	4.6	0.581	0.35	3.85±0.11	3.69±0.14	3.76±0.08
MegaTTS-Continue	580	5.8	0.615	0.39	3.93±0.08	3.85±0.09	3.89±0.09
MobileSpeech-Continue	580	3.1	0.688	0.09	4.06±0.07	4.02±0.08	4.05±0.10

Table 1: The results of different zero-shot TTS models on the LibriSpeech test-clean dataset.

Model	CMOS-Q ↑	CMOS-P ↑	CMOS-S ↑
GT codec	+0.07	+0.15	+0.03
MegaTTS 2	-0.16	-0.05	-0.24
MobileSpeech	0.00	0.00	0.00

Table 2: The CMOS (comparative Mean Opinion Score) results of different models on the Chinese Mandarin test set.

tors. We randomly choose 50 samples from the test set of each dataset for the subjective evaluation and each audio is listened to by at least 20 testers. We analyze the MOS in three aspects: MOS-Q (Quality: clarity, high-frequency, and original timbre reconstruction), MOS-P (Prosody: naturalness of pitch, energy, and duration), and MOS-S (Speaker similarity). We also analyze the CMOS in terms of audio quality and speech prosody. See Appendix E for details.

4.2 LibriSpeech Evaluation

Regarding the inference phase, to ensure fair comparisons, we followed the experimental protocols outlined in VALL-E (Wang et al., 2023) and employed the LibriSpeech test-clean dataset (Panayotov et al., 2015), ensuring no overlap with our training data. We specifically utilized samples from the LibriSpeech test-clean dataset with durations ranging from 4 to 10 seconds, resulting in a subset of 2.2 hours. Following VALL-E approach, we use the whole transcription and the first 3 seconds of the utterance as the phoneme and acoustic prompts respectively, and ask the model to generate the continuations, the continue setting of Table 1 is referred to this.

By observing Table 1, the following conclusions can be drawn: (1) Under the same scale of 580 hours training data of libriTTS (where yourtts is trained on a combined dataset of VCTK, LibriTTS, and TTS-Portuguese (Casanova et al., 2022a)), MobileSpeech significantly outperforms YourTTS

in terms of SPK and MOS, surpassing the current SOTA models NaturalSpeech2 and MegaTTS. This indicates that MobileSpeech is capable of generating higher audio quality, better prosody, and more similar audio. The outstanding results are attributed to the utilization of the baseline generated by the MASK parallel codec and the additional design of the SMD module, Speaker Prompt module in MobileSpeech. (2) In terms of audio quality, MobileSpeech also achieves comparable results to VALL-E, which was trained on a large-scale dataset of 60,000 hours, with a speaker similarity metric difference of only 0.046. This difference is attributed to the disparity in the training data volume. (3) MobileSpeech achieves the best results in terms of robustness measured by the Wer metric, even surpassing VALL-E by 0.7. This is because the autoregressive structure of VALL-E often suffers from cumulative errors, resulting in lower robustness. Our experiments reveal that increasing the amount of data can alleviate the issue of high robustness. (4) MobileSpeech exhibits remarkable performance in terms of the Real-Time Factor (RTF) metric, as it can generate 10 seconds of audio within one second. Compared to VALL-E, MobileSpeech achieves an 11-fold speed improvement, and compared to MegaTTS and NaturalSpeech2, it achieves a 4-fold speed improvement. Although YourTTS achieves an RTF of 0.22, its simple structural design leads to unsatisfactory performance in terms of audio quality, speaker similarity, and model robustness.

4.3 Mandarin Evaluation

As our inference dataset for Mandarin Chinese, we compiled a collection of 200 real-world mobile-end recordings with arbitrary durations ranging from 3 to 10 seconds, representing a mixture of clean and noisy scenarios. This dataset serves as our Mandarin Chinese test set. We employed CMOS to measure the performance of MobileSpeech and the

Model	WER ↓	SPK ↑
MobileSpeech w/o durprompt	3.6	0.638
MobileSpeech w/o onechannel	4.2	0.614
MobileSpeech	3.1	0.688

Table 3: The ablation experiments of the SMD module and the Speaker Prompt module were conducted to test SPK and WER.

Model	MOS-Q ↑	MOS-P ↑	MOS-S ↑
MobileSpeech w/o durprompt	4.03±0.05	3.90±0.11	3.97±0.12
MobileSpeech w/o onechannel	3.98±0.12	3.95±0.09	3.92±0.11
MobileSpeech	4.06±0.07	4.02±0.08	4.05±0.10

Table 4: The ablation experiments of the SMD module and the Speaker Prompt module were conducted to test MOS.

current SOTA model MegaTTS2 on the Mandarin Chinese dataset. Considering that MegaTTS2 is a paid model, we incurred significant costs to conduct this experiment.

The experimental results, as shown in Table 2, demonstrate that MobileSpeech outperforms MegaTTS2 comprehensively in terms of MOS-Q, MOS-P, and MOS-S under various conditions, including large-scale training corpora, different language categories, varying prompt lengths, and noisy speech prompts. Particularly, MobileSpeech achieves a significantly higher score of 0.24 in terms of speaker similarity compared to MegaTTS2. By conducting a cross-comparison between MobileSpeech and GT Codec in terms of CMOS-Q, CMOS-P, and CMOS-S metrics, we can observe that MobileSpeech exhibits superior performance in terms of speaker similarity but slightly lags behind in audio prosody.

4.4 Parameters Evaluation

To begin, we present a comprehensive account of the parameters for all state-of-the-art models. In the VALL-E framework, the AR model comprises 229 million parameters, while the NAR model consists of 235 million parameters, resulting in a total of 465 million parameters. In the NaturalSpeech2 model, the overall parameter count is 435 million. Within the MegaTTS system, the VQ-GAN model encompasses 243 million parameters, the LLM model comprises 124 million parameters, and the hifigan vocoder consists of 106 million parameters, resulting in a cumulative parameter count of 473 million. As outlined in the accompanying figure, the MobileSpeech model has a total parameter count of 207 million (detailed parameters can be

iterations	RTF ↓	WER ↓	SPK ↑
24	0.20	3.1	0.698
16	0.15	3.0	0.696
8	0.09	3.1	0.688
4	0.07	4.2	0.615
1	0.05	5.7	0.483

Table 5: The impact of different iterations on the first channel of MobileSpeech.

found in Appendix C of the paper), demonstrating its status as the most lightweight option.

4.5 Ablation experiment

We conducted ablation experiments on the LibriSpeech test-clean dataset. Firstly, we omitted fine-grained prompt duration and directly used prompt acoustic tokens to guide the generation of target text. Additionally, in the SMD module, prompt acoustic tokens were solely used as the input of SMD. This version was referred to as MobileSpeech w/o durprompt. For the MobileSpeech w/o onechannel version, we replaced the SMD module with randomly selecting a layer for mask prediction. The results are presented in Tables 3 and 4. It can be observed that both MobileSpeech w/o durprompt and MobileSpeech w/o onechannel exhibited certain decreases in WER, SPK, and MOS values compared to MobileSpeech. This suggests that each component of MobileSpeech’s design is necessary. During our experimental process, we found that the MobileSpeech w/o durprompt version exhibited a certain amount of disfluency, which affected the overall prosody of the audio. This indicates that fine-grained duration prompts indeed contribute to better modeling of duration information.

We also conducted ablation experiments on the number of iterations in the first channel of MobileSpeech, which has a significant impact on the inference speed. As shown in Table 5, we observed that the number of iterations is inversely proportional to the real-time factor (RTF), and blindly increasing the number of iterations does not necessarily improve the model’s performance. However, when the number of iterations is reduced to 1, the model exhibits significant degradation in its capabilities. Consequently, we selected 8 iterations as the final choice, striking a balance between performance and efficiency.

5 Conclusion

In this paper, we propose MobileSpeech, a Fast and Lightweight Framework for Mobile zero-shot Text-to-Speech. MobileSpeech achieves a unique combination of audio naturalness and similarity performance, comparable to other zero-shot TTS systems on English corpora, while significantly enhancing inference speed through its distinctive SMD and Speaker modules. MobileSpeech stands as the first real-time zero-shot TTS system deployable at the edge, demonstrating state-of-the-art performance on Chinese corpora through extensive training on large-scale Chinese language data. We not only validate the significance of each module in MobileSpeech through ablation experiments but also analyze the effects of different configurations and iterations to adapt to various deployment scenarios in mobile environments.

6 Acknowledgments

This work was supported in part by the National Key R&D Program of China under Grant No.2022ZD0162000

7 Limitations and Future Work

In this section, we will analyze the limitations of MobileSpeech and discuss potential future work.

Failure cases During our experiments in real-world scenarios (Figure 2), we observed occasional instances of stuttering and unstable pitch in MobileSpeech. This could be attributed to noisy input prompts or the possibility that the mask mechanism excessively focuses on local positions. In future work, we aim to enhance the model’s noise robustness and design a more robust mask mechanism that considers a broader context.

New Tasks: The current zero-shot TTS task itself has significant potential for expansion. Most existing models primarily focus on similarity in timbre. However, future research can extend zero-shot TTS to include similarity in rhythm, emotion, and language. It could even involve using multiple audio segments to control different parts or combining with text style control, representing meaningful new tasks.

Broader impacts: Since MobileSpeech could synthesize speech that maintains speaker identity, it may carry potential risks in misuse of the model, such as spoofing voice identification or impersonating a specific speaker. To mitigate such risks, it is possible to build a detection model to discriminate

whether an audio clip was synthesized by MobileSpeech.

References

- Sercan Arik, Jitong Chen, Kainan Peng, Wei Ping, and Yanqi Zhou. 2018. Neural voice cloning with a few samples. *Advances in neural information processing systems*, 31.
- Zalán Borsos, Raphaël Marinier, Damien Vincent, Eugene Kharitonov, Olivier Pietquin, Matt Sharifi, Dominik Roblek, Olivier Teboul, David Grangier, Marco Tagliasacchi, et al. 2023a. Audiolm: a language modeling approach to audio generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Zalán Borsos, Matt Sharifi, Damien Vincent, Eugene Kharitonov, Neil Zeghidour, and Marco Tagliasacchi. 2023b. Soundstorm: Efficient parallel audio generation. *arXiv preprint arXiv:2305.09636*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Edresson Casanova, Arnaldo Candido Junior, Christopher Shulby, Frederico Santos de Oliveira, João Paulo Teixeira, Moacir Antonelli Ponti, and Sandra Aluísio. 2022a. Tts-portuguese corpus: a corpus for speech synthesis in brazilian portuguese. *Language Resources and Evaluation*, 56(3):1043–1055.
- Edresson Casanova, Julian Weber, Christopher D Shulby, Arnaldo Candido Junior, Eren Gölge, and Moacir A Ponti. 2022b. Yourtts: Towards zero-shot multi-speaker tts and zero-shot voice conversion for everyone. In *International Conference on Machine Learning*, pages 2709–2720. PMLR.
- Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. 2022. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11315–11325.
- Mingjian Chen, Xu Tan, Bohan Li, Yanqing Liu, Tao Qin, Sheng Zhao, and Tie-Yan Liu. 2021. Adaspeech: Adaptive text to speech for custom voice. *arXiv preprint arXiv:2103.00993*.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518.
- Alexandre Défossez, Jade Copet, Gabriel Synnaeve, and Yossi Adi. 2022. High fidelity neural audio compression. *arXiv preprint arXiv:2210.13438*.

- Chenpeng Du, Yiwei Guo, Feiyu Shen, Zhijun Liu, Zheng Liang, Xie Chen, Shuai Wang, Hui Zhang, and Kai Yu. 2023. Unicats: A unified context-aware text-to-speech framework with contextual vq-diffusion and vocoding. *arXiv preprint arXiv:2306.07547*.
- Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. 2019. Mask-predict: Parallel decoding of conditional masked language models. *arXiv preprint arXiv:1904.09324*.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al. 2020. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhota, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460.
- R Huang, MWY Lam, J Wang, D Su, D Yu, Y Ren, and Z Zhao. 2022a. Fastdiff: A fast conditional diffusion model for high-quality speech synthesis. In *IJCAI International Joint Conference on Artificial Intelligence*, pages 4157–4163. IJCAI: International Joint Conferences on Artificial Intelligence Organization.
- Rongjie Huang, Chunlei Zhang, Yongqi Wang, Dongchao Yang, Luping Liu, Zhenhui Ye, Ziyue Jiang, Chao Weng, Zhou Zhao, and Dong Yu. 2023. Make-a-voice: Unified voice synthesis with discrete representation. *arXiv preprint arXiv:2305.19269*.
- Sung-Feng Huang, Chyi-Jiunn Lin, Da-Rong Liu, Yi-Chen Chen, and Hung-yi Lee. 2022b. Meta-tts: Meta-learning for few-shot speaker adaptive text-to-speech. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:1558–1571.
- Shengpeng Ji, Jialong Zuo, Minghui Fang, Ziyue Jiang, Feiyang Chen, Xinyu Duan, Baoxing Huai, and Zhou Zhao. 2023. Textrolspeech: A text style control speech corpus with codec language text-to-speech models. *arXiv preprint arXiv:2308.14430*.
- Ziyue Jiang, Jinglin Liu, Yi Ren, Jinzheng He, Chen Zhang, Zhenhui Ye, Pengfei Wei, Chunfeng Wang, Xiang Yin, Zejun Ma, et al. 2023a. Mega-tts 2: Zero-shot text-to-speech with arbitrary length speech prompts. *arXiv preprint arXiv:2307.07218*.
- Ziyue Jiang, Yi Ren, Zhenhui Ye, Jinglin Liu, Chen Zhang, Qian Yang, Shengpeng Ji, Rongjie Huang, Chunfeng Wang, Xiang Yin, et al. 2023b. Mega-tts: Zero-shot text-to-speech at scale with intrinsic inductive bias. *arXiv preprint arXiv:2306.03509*.
- Jacob Kahn, Morgane Rivière, Weiyi Zheng, Evgeny Kharitonov, Qiantong Xu, Pierre-Emmanuel Mazaré, Julien Karadayi, Vitaliy Liptchinsky, Ronan Collobert, Christian Fuegen, et al. 2020. Libri-light: A benchmark for asr with limited or no supervision. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7669–7673. IEEE.
- Minki Kang, Dongchan Min, and Sung Ju Hwang. 2022. Any-speaker adaptive text-to-speech synthesis with diffusion models. *arXiv preprint arXiv:2211.09383*.
- Eugene Kharitonov, Damien Vincent, Zalán Borsos, Raphaël Marinier, Sertan Girgin, Olivier Pietquin, Matt Sharifi, Marco Tagliasacchi, and Neil Zeghidour. 2023. Speak, read and prompt: High-fidelity text-to-speech with minimal supervision. *arXiv preprint arXiv:2302.03540*.
- Jaehyeon Kim, Sungwon Kim, Jungil Kong, and Sungho Yoon. 2020. Glow-tts: A generative flow for text-to-speech via monotonic alignment search. *Advances in Neural Information Processing Systems*, 33:8067–8077.
- Jaehyeon Kim, Jungil Kong, and Juhee Son. 2021a. Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech. In *International Conference on Machine Learning*, pages 5530–5540. PMLR.
- Jaehyeon Kim, Jungil Kong, and Juhee Son. 2021b. Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech. In *International Conference on Machine Learning*, pages 5530–5540. PMLR.
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. *Advances in Neural Information Processing Systems*, 33:17022–17033.
- Felix Kreuk, Gabriel Synnaeve, Adam Polyak, Uriel Singer, Alexandre Défossez, Jade Copet, Devi Parikh, Yaniv Taigman, and Yossi Adi. 2022. Audiogen: Textually guided audio generation. *arXiv preprint arXiv:2209.15352*.
- Matthew Le, Apoorv Vyas, Bowen Shi, Brian Karrer, Leda Sari, Rashel Moritz, Mary Williamson, Vimal Manohar, Yossi Adi, Jay Mahadeokar, et al. 2023. Voicebox: Text-guided multilingual universal speech generation at scale. *arXiv preprint arXiv:2306.15687*.
- Yoonhyung Lee, Joongbo Shin, and Kyomin Jung. 2022. Bidirectional variational inference for non-autoregressive text-to-speech. In *International conference on learning representations*.
- Yinghao Aaron Li, Cong Han, Vinay S Raghavan, Gavin Mischler, and Nima Mesgarani. 2023. Styletts 2: Towards human-level text-to-speech through style diffusion and adversarial training with large speech language models. *arXiv preprint arXiv:2306.07691*.

- Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. 2017. Montreal forced aligner: Trainable text-speech alignment using kaldi. In *Interspeech*, volume 2017, pages 498–502.
- Chenfeng Miao, Shuang Liang, Minchuan Chen, Jun Ma, Shaojun Wang, and Jing Xiao. 2020. Flow-tts: A non-autoregressive network for text to speech based on flow. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7209–7213. IEEE.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5206–5210. IEEE.
- Yi Ren, Chenxu Hu, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2020. Fastspeech 2: Fast and high-quality end-to-end text to speech. In *International Conference on Learning Representations*.
- Yi Ren, Jinglin Liu, and Zhou Zhao. 2021. Portaspeech: Portable and high-quality generative text-to-speech. *Advances in Neural Information Processing Systems*, 34:13963–13974.
- Yi Ren, Yangjun Ruan, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2019. Fastspeech: Fast, robust and controllable text to speech. *Advances in neural information processing systems*, 32.
- Kai Shen, Zeqian Ju, Xu Tan, Yanqing Liu, Yichong Leng, Lei He, Tao Qin, Sheng Zhao, and Jiang Bian. 2023. Naturalspeech 2: Latent diffusion models are natural and zero-shot speech and singing synthesizers. *arXiv preprint arXiv:2304.09116*.
- Hubert Siuzdak. 2023. Vocos: Closing the gap between time-domain and fourier-based neural vocoders for high-quality audio synthesis. *arXiv preprint arXiv:2306.00814*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang, Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. 2022. Phenaki: Variable length video generation from open domain textual description. *arXiv preprint arXiv:2210.02399*.
- Chengyi Wang, Sanyuan Chen, Yu Wu, Ziqiang Zhang, Long Zhou, Shujie Liu, Zhuo Chen, Yanqing Liu, Huaming Wang, Jinyu Li, et al. 2023. Neural codec language models are zero-shot text to speech synthesizers. *arXiv preprint arXiv:2301.02111*.
- Yuxuan Wang, Daisy Stanton, Yu Zhang, RJ-Skerry Ryan, Eric Battenberg, Joel Shor, Ying Xiao, Ye Jia, Fei Ren, and Rif A Saurous. 2018. Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis. In *International conference on machine learning*, pages 5180–5189. PMLR.
- Dongchao Yang, Songxiang Liu, Rongjie Huang, Jinchuan Tian, Chao Weng, and Yuexian Zou. 2023a. Hifi-codec: Group-residual vector quantization for high fidelity audio codec. *arXiv preprint arXiv:2305.02765*.
- Dongchao Yang, Jinchuan Tian, Xu Tan, Rongjie Huang, Songxiang Liu, Xuankai Chang, Jiatong Shi, Sheng Zhao, Jiang Bian, Xixin Wu, et al. 2023b. Uniaudio: An audio foundation model toward universal audio generation. *arXiv preprint arXiv:2310.00704*.
- Neil Zeghidour, Alejandro Luebs, Ahmed Omran, Jan Skoglund, and Marco Tagliasacchi. 2021. Soundstream: An end-to-end neural audio codec. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:495–507.
- Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu. 2019. Libritts: A corpus derived from librispeech for text-to-speech. *Interspeech 2019*.
- Ziqiang Zhang, Long Zhou, Chengyi Wang, Sanyuan Chen, Yu Wu, Shujie Liu, Zhuo Chen, Yanqing Liu, Huaming Wang, Jinyu Li, et al. 2023. Speak foreign languages with your own voice: Cross-lingual neural codec language modeling. *arXiv preprint arXiv:2303.03926*.

A Discrete Acoustic Codec

Since audio is typically stored as a sequence of 16-bit integer values, a generative model is required to output $2^{16} = 65536$ probabilities per timestep to synthesize the raw audio. In addition, the audio sample rate exceeding ten thousand leads to an extraordinarily long sequence length, making it more intractable for raw audio synthesis. In these days, neural acoustic codec (Zeghidour et al., 2021; Défossez et al., 2022; Yang et al., 2023a; Siuzdak, 2023), with their ability of reconstructing high-quality audio at very low bitrates, sub-sequently allowed for extending discrete modeling to audio signals as diverse as multi-speaker speech. To be specific, Soundsteam (Zeghidour et al., 2021) which relies on a model architecture composed by a fully convolutional encoder/decoder network and a residual vector quantizer (RVQ) can efficiently compress speech. Encodec (Défossez et al., 2022) consists in a streaming encoder-decoder architecture with quantized latent space trained in an end-to-end fashion. Hificodec (Yang et al., 2023a) propose a group-residual vector quantization (GRVQ)

technique To reduce the number of quantizers. Vocos (Siuzdak, 2023) closing the gap between time-domain and fourier-based neural vocoders for high-quality audio synthesis. After our detailed experiments, MobileSpeech finally selected encdec and vocos as the encoder and decoder of acoustic token.

B Codec Modeling

In this section, we provide a detailed comparison of different modeling approaches for the RVQ structure based on discrete codecs. As depicted in Figure 3, (a) represents a fully sequential unfolding of each frame’s channels followed by autoregressive modeling, which incurs significant time overhead as the number of channels increases. (b) involves splitting the codec generation into two stages, where the first stage autoregressively generates the codec for the first channel, and the second stage concurrently generates the codecs for the remaining channels. (c) adopts global attention in the temporal dimension and local attention in the channel dimension for modeling. All three methods (a), (b), and (c) are fundamentally autoregressive models, but due to limitations in inference speed and memory consumption, they are not suitable for deployment on mobile devices. In contrast, MobileSpeech adopts a mask-based parallel generation approach, offering a potential deployment solution.

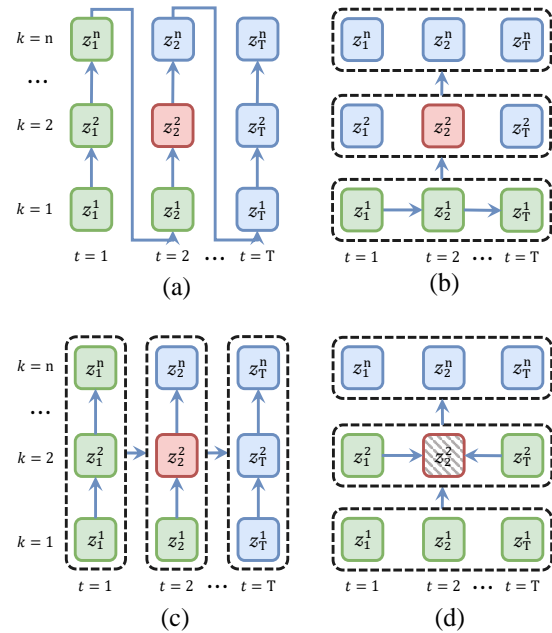


Figure 3: Based on the RVQ structure, different articles adopt distinct modeling approaches for discrete codecs, where (a) SpearTTS (Kharitonov et al., 2023), (b) VALL-E (Wang et al., 2023), (c) Uniaudio (Yang et al., 2023b) and (d) MobileSpeech.

C Model Structure

In this section, we introduce the detailed structure of proposed MobileSpeech. To be specific, we will introduce the architecture of text encoder, prompt duration encoder, prompt encoder, cross-attention and ConFormer (Gulati et al., 2020) used in MobileSpeech. Text encoder and prompt duration encoder share the same architecture which is comprised of several FFT blocks as used by FastSpeech2 (Ren et al., 2020). The value of probability parameter α in MobileSpeech is 0.6. The detailed configuration information is shown in the table 6.

D Algorithm

We provide a detailed description of the P_{rank} algorithm in the Algorithm 1 provided below.

E MOS Evaluation

Below, we provide a detailed explanation of the requirements and standards we employed for measuring Mean Opinion Score (MOS).

MOS50_mainresult

Instructions:

- For each text, you will hear the example audio and seven audio clips a,b,c,d,e,f,g. Please rate the seven audio recordings a,b,c,d,e,f,g.
- Please analyze the MOS from three aspects (MOS-Q/MOS-P/MOS-S), considering each indicator individually without considering the impact of the other.
- MOS-Q (Quality: clarity, high-frequency, and original timbre reconstruction), MOS-P (Prosody: naturalness of pitch, energy, and duration), and MOS-S (Speaker similarity). MOS-S measures the similarity between each piece of audio and the sample audio.
- Please carefully listen to the provided examples, assigning a score ranging from 5.0 to 1.0 to each example.
- If you provide a score without listening to the examples, your response will not be accepted. Additionally, we have restricted the scoring time for each audio segment.
- For optimal results, please wear headphones and choose a quiet environment.

The MOS50_mainresult can be found at <https://drive.google.com/file/d/17bW->

7. Name:

testid07_librispeech_61-70970-0000

Text:

YOUNG FITZTOOTH HAD BEEN COMMANDED TO HIS MOTHER'S CHAMBER SO SOON AS HE HAD COME OUT FROM HIS CONVERSE WITH THE SQUIRE

A	MOS-Q: <input type="text"/>	MOS-P: <input type="text"/>	MOS-S: <input type="text"/>
B	MOS-Q: <input type="text"/>	MOS-P: <input type="text"/>	MOS-S: <input type="text"/>
C	MOS-Q: <input type="text"/>	MOS-P: <input type="text"/>	MOS-S: <input type="text"/>
D	MOS-Q: <input type="text"/>	MOS-P: <input type="text"/>	MOS-S: <input type="text"/>
E	MOS-Q: <input type="text"/>	MOS-P: <input type="text"/>	MOS-S: <input type="text"/>
F	MOS-Q: <input type="text"/>	MOS-P: <input type="text"/>	MOS-S: <input type="text"/>
G	MOS-Q: <input type="text"/>	MOS-P: <input type="text"/>	MOS-S: <input type="text"/>

Figure 4: MOS evaluation procedure

Module	Hyper-parameters	Parameters Number
Prompt Text Encoder	Number of FFT Blocks	2
	Conv1D kernel size	3
	Conv1D input channels	512
	Attention heads	8
Text Encoder	Number of FFT Blocks	4
	Conv1D kernel size	3
	Conv1D input channels	512
	Attention heads	8
Prompt Encoder	Number of layers	2
	Hidden channels	512
	Conv1D kernel size	3
	Filter channels	2048
	Attention heads	4
Duration Extractor	Number of attention layers	10
	Hidden channels	512
	Attention heads	8
Duration Predictor	Number of attention layers	10
	Hidden channels	512
	Attention heads	8
Prompt Duration Encoder	Number of FFT Blocks	2
	Conv1D kernel size	3
	Conv1D input channels	512
	Attention heads	8
Conformer	Number of Encoder layers	16
	Attention dim	512
	Attention heads	16
	Linear units	1024
	Positionwise Conv1D Kernel	5
	Dropout probability	0.1

Table 6: The configuration of modules of MobileSpeech.

Algorithm 1 P_{rank} algorithm

Require: The channels set $\{c_i\}_{i=2}^N$ and decreasing corresponding weights $\{w_i\}_{i=2}^N$; where $w_j > w_{j+1}, j \in \{2, 3, \dots, N-1\}$

Initialize $\mathcal{C}_U = \emptyset$;

for each c_i **do**

Repeat c_i for w_i times;

Add repeated c_i to \mathcal{C}_U ;

end for

Random sample a channel c from \mathcal{C}_U ;
